

Design and Verification of the Second-Generation Precipitation Radar Processor/Controller

Mark A. Fischman, Andrew C. Berkun, William W. Chun, and Eastwood Im
Jet Propulsion Laboratory, California Institute of Technology, Radar Science and Engineering Section
M/S 300-241, 4800 Oak Grove Dr, Pasadena, CA 91109-8099, USA
Tel: (818) 354-3862, Fax: (818) 393-6943, e-mail: mafisch@mail1.jpl.nasa.gov

Ray Andraka
Andraka Consulting Group, Inc.
North Kingstown, RI 02852-1666
ray@andraka.com

Abstract—NASA/JPL has developed the on-board processing electronics for the Second-Generation Precipitation Radar (PR-2)—a follow-on to the 14 and 36 GHz precipitation radar aboard the Global Precipitation Measurement (GPM) mission—to support advanced capabilities such as adaptive scanning and compression of rainfall science data. We describe the design and experimental performance of the spaceborne PR-2 processor/controller module, which is based upon field-programmable gate array (FPGA) technology. Test results demonstrate a high timing efficiency of about 90% for the adaptive scan and very low range side-lobes of -75 dB for the radar's receive processor.

I. INTRODUCTION

Advanced radar technology is sought in the upcoming Global Precipitation Measurement (GPM) constellation of earth observing satellites in order to measure the detailed structure and distribution of rainfall. One instrument currently being developed is the Dual-Frequency Precipitation Radar (DPR), a Ku- and Ka-band (14 and 36 GHz) sensor from the National Space Development Agency of Japan. DPR uses a conventional phased-array / pulsed radar technology that has been inherited from the Tropical Rainfall Measuring Mission (TRMM) [1]. While DPR is the first satellite instrument to offer higher rainfall detection sensitivity through its 36 GHz channel, the key observation capabilities of the radar—the 5 km horizontal resolution and 250 km swath coverage—are essentially the same as for the original TRMM radar.

To truly make the gains in spatial resolution and swath width that will be needed to observe rain down to the scale of convective cells, NASA/JPL is leading the development of some breakthrough technologies for a Second-Generation Precipitation Radar (PR-2), as a follow-on to GPM. In this paper, we will cover our group's recent progress in developing and testing one particular aspect of the PR-2 design—the electronic on-board processor/controller module. The design approach is to use high-density, radiation-tolerant field-programmable gate arrays (FPGAs) to realize greater compression of science data

(through adaptive scanning of the rain target) and higher fidelity imaging of the rain profiles (through very high-throughput processing of the received radar echoes).

II. ADAPTIVELY SCANNING RADAR

A. Timing Solution

Because rain tends to be sparsely distributed over large scales (> 100 km), there has been an interest in developing adaptive scanning radar systems that have the built-in intelligence to auto-target areas of rain [2]. Adaptive scanning becomes all the more necessary as the observation requirements for a Second-Generation Precipitation Radar move toward high horizontal footprint resolutions of 2 km and large swath widths extending up to 500 km. Compared to TRMM or to the precipitation radar observations planned for the GPM mission, the PR-2 represents an 8-fold increase in the number of radar beam locations per swath area (a 4-fold increase due to resolution doubling and a 2-fold increase in swath width). Given the limited dwell time of a 7 km/s low-earth orbiting satellite, the timing sequence for interlacing the transmit and receive beams becomes especially critical. The radar must be able to selectively flag only those beam locations that have scientifically valuable rainfall data and then generate a transmit/receive timing solution that can capture the $N \geq 60$ or so independent radar looks needed to accurately estimate the reflectivity of each rain bin.

We have developed a unique timing solution for the electronically scanning PR-2 radar that can auto-target rain having any possible spatial distribution. The scan sequence consists of: a *locator sweep* over the entire cross-track swath to roughly estimate the rainfall distribution; a *bubble-sort* algorithm to rank beam locations from highest to lowest rainfall; and a *high-resolution* sweep in which the cross-track beam positions with the greatest rainfall are observed over a much longer integration period. The timing solution is unlike the predetermined, left-to-right scan sequences used in conventional radars such as TRMM that may not make the most efficient use of the available dwell time. The algorithm for the PR-2 begins with, as an input, a desired sequence of beam locations to be observed for either locator or high-

resolution sweeps. The timing algorithm then can be posed as a counting problem: In what order should the beams be counted, and how should the pulse-repetition interval (PRI) be varied, to gather the largest possible number of independent radar looks from the rain scene? Simultaneously, how can it be guaranteed that there are no collisions between the radar's transmit and receive intervals? Our timing solution is based on several rules to meet these requirements:

- The receive window length (the slant altitude range for the rain scene) is held constant throughout the sweep.
- The beam sequence always moves from nadir to alternating left and right sides of the cross-track scan—that is, from the shortest to longest range delays to the target—so that the average PRI is monotonically increasing.
- The number of echoes-in-flight (EIFs) is determined by the largest integer number of PRIs that can fit within the shortest (first) range delay of the scan sequence; this EIF value is held constant throughout the scan.
- The beam scanning order is counted in an alternating pattern (e.g., beam numbers 1,2,1,2,3,4,3,4,...) so that the satellite has time to move laterally between echoes. This ensures that multiple radar looks of the same rain bin are statistically independent.

As will be discussed next, the above rules can be realized in digital circuitry to make a robust adaptive scanning controller for the PR-2.

B. Control and Timing Unit (CTU)

The adaptive scanning solution for the PR-2 has been implemented onto a single Xilinx Virtex series FPGA, whose design has been described in detail in our previous work [3]. This Control and Timing Unit (CTU) FPGA generates the transmit and receive timing for the radar hardware and commands the PR-2's phased-array antenna for electronic steering of the beam. The CTU firmware, written in the Verilog hardware description language, has been designed with a custom state machine architecture, rather than a microprocessor core, in order to respond with zero latency to the large number of timing interrupts (up to 15–20) associated with each echo-in-flight (EIF).

Over the last year, the CTU design has been upgraded to sort rain reflectivities and adaptively scan over the full 500 km, 250 beam swath. Newly written modules include the logic for calculating the optimum number of EIFs for an arbitrary scan sequence, a floating-point accumulator for averaging locator sweep echoes, and a driver for storing the radar echoes to a solid-state data recorder. The CTU design currently synthesizes to an equivalent gate count of 79,000 logic gates.

The CTU was tested over entire adaptive scan cycles to measure the *timing efficiency* of the radar—that is, the percentage of the total sweep time dedicated to transmitting a pulse or receiving an echo. Verilog simulations were run with several types of rain profiles to test varying levels of difficulty for the auto-targeting algorithm: 1) a *best case* profile where all heavy rain is

rain profile	total scan time (ms)	timing efficiency (%)
worst case	277.5	82.9
random	260.0	88.5
best case	245.3	93.7

Table 1: Timing efficiency results for the adaptive scan algorithm.

located near the nadir direction (i.e., very little range delay variation); 2) a *worst case* profile in which there is a small amount of rain at nadir and the rest of the rain is at the extreme left and right edges of the swath (a large range delay variation between beams); and 3) a *random case* where the rainfall is uniformly distributed over the swath. Table 1 lists the timing efficiency results for the full 250 beam scan. The results indicate that our timing solution for the adaptive scan will range from 83–94%, depending on how the rain is distributed. (Efficiencies of 90% or above are considered excellent.)

The CTU tests described above were run with about 10% of the beams (24 out of 250) selected for the high-resolution sweep. It should be pointed out that this represents just one of many possible settings for the adaptive scan. Numeric registers were designed into the FPGA so that the timing solution parameters could easily be changed through software. Currently, a time window of 170 ms out of 260 ms is available for the high resolution sweep. Given a PRI of 200 μ s or more for rainfall profiling, this allows time to sample approximately 800 echoes in hi-res mode. The number of independent echoes was therefore set to 32 so that the product (24 beams \times 32 looks = 768 echoes) can fit within the time limit. For comparison, consider a conventional scanning radar with the same total available time of 260 ms per cross-track scan. The number of independent echoes that can be measured is then: total scan time / (total number of beams \times PRI) \approx 5 echoes. It follows that in our adaptive scan configuration tested with 10% auto-targeting coverage, the number of independent radar looks improves by a factor of $32/5 = 6.4$ compared to a traditional radar. Because of the built-in flexibility of the CTU software registers, it is easy for the PR-2 operation to be adjusted to make tradeoffs between swath coverage versus number of radar looks, depending on the type of weather being observed and on requirements for the rain rate retrieval algorithms.

III. DATA PROCESSOR

The PR-2 operates as a linear FM chirp radar which must compress the received echoes into complex-valued pulses with respect to range. Our team has developed and implemented this pulse-compression operation onto two Virtex-1000 Data Processor FPGAs using a bit-serial architecture [3]. Together the two FPGA parts simultaneously process four receive channels (Ku- and Ka-band, co- and cross-polarizations) at a peak throughput rate of 40×10^9 ops/s. The technique was proven successful aboard an airborne prototype of the PR-2, last deployed in January 2003 for the AMSR-E (Advanced Microwave Scanning Radiometer) validation experiment in Japan. In this campaign, the radar collected the first comprehensive set of precipitation data

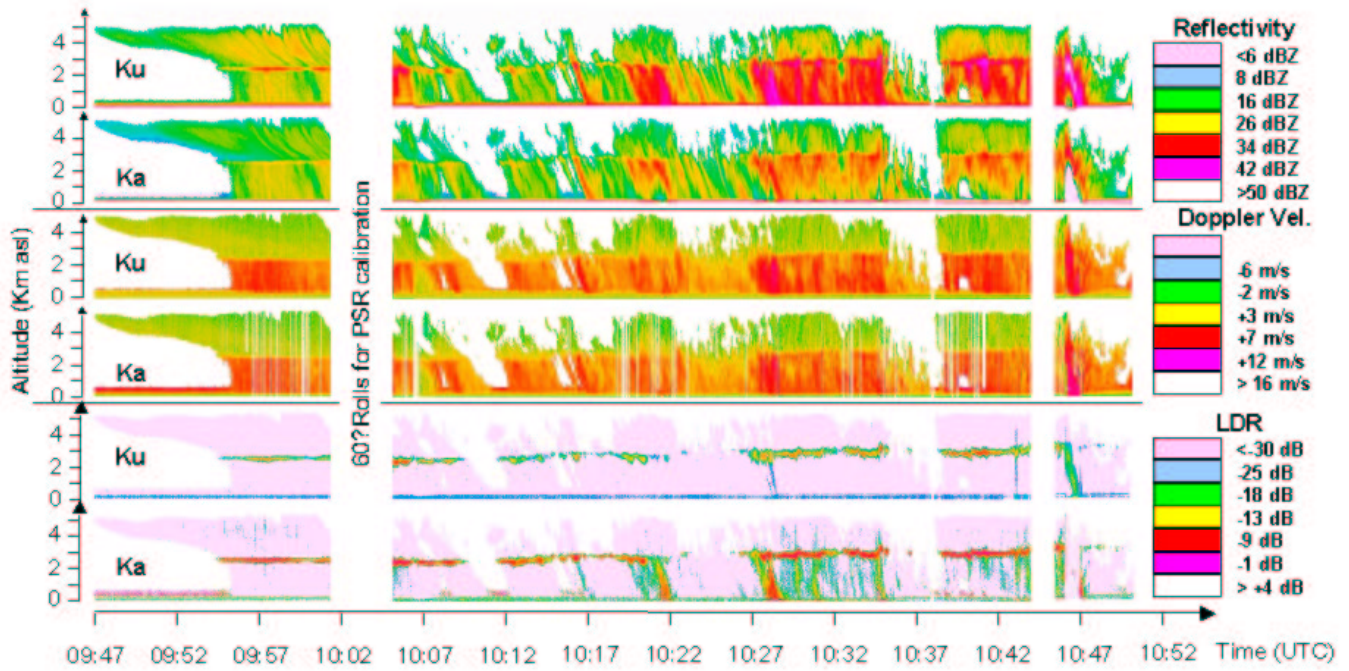


Figure 1: Multi-parameter, dual-frequency (14 and 36 GHz) rainfall data collected by the airborne version of the PR-2 FPGA processors. This 500 km flight line was flown on the NASA P-3 Orion aircraft on January 23, 2003, over the Pacific Ocean, east of the main island of Japan, as part of the AMSR-E validation experiment. Graphs from top to bottom show the along-track (nadir) radar observations for: co-polarized rain reflectivity data, vertical Doppler velocity measurements, and linear depolarization ratios.

at both 14 and 36 GHz bands (see Fig. 1).

New design features were added to the spaceborne version of the Data Processor to accommodate adaptive scanning. The timing inputs were revised to handle the averaging of multiple echoes in memory with a staggered beam sequence. Also, the fixed-point logic for averaging echoes has been replaced with a 32-bit floating-point accumulator to increase the processor's dynamic range.

IV. TEST RESULTS FOR SPACEBORNE HARDWARE

The physical layout for the PR-2 processor/controller module is composed of a pair of multi-layer printed circuit boards and an aluminum chassis (see Fig. 2), both which have been specially designed for a space environment. Two main aspects of the design that increase the reliability and survivability of the FPGA hardware in space are: 1) configuration memory scrubbing to recover from radiation-induced logic upsets; and 2) thermal design of the chassis for conductive cooling of the Data Processor FPGAs. The processor's operation has been successfully confirmed in a series of thermal chamber tests over an ambient temperature range of -20 to $+70$ °C.

Fig. 3 shows the experiment setup for the PR-2 processor, along with a PC-based test fixture to control and acquire data from the prototype boards. The PC computer interfaces to the boards via a 32-channel National Instruments Data I/O (DIO) card. Digital commands are sent from the DIO card to an on-board Actel bus controller FPGA, programmed to convert these




Figure 2: Space-grade chassis with assembled printed circuit boards for the PR-2 processor/controller. Machined heat sink blocks for conductive cooling of the FPGA ball-grid array packages [inset].




Figure 3: Test setup for spaceborne PR-2 chassis. (a) Benchtop testing with boards interfaced to a PC computer with LabView control software. (b) Processor boards + radar echo signal generator in the thermal chamber.

commands into local bus transactions for the Virtex FPGAs. We developed a graphical user interface (GUI) in the LabView visual language to provide the various controls and indicators needed for benchtop operation of the PR-2 processor. With the benchtop test fixture, entire adaptive scan cycles can be run in real time and memory contents examined to verify the radar processor operation.

A. Pulse-Compressed Output

The PR-2 processor boards were tested in the thermal chamber in free-running mode with all four ADC input channels connected to a custom-made signal generator. This signal generator waveform is programmable through read-only memory to simulate the chirp radar response to a rain target at offset video frequencies (after digitization of the received signal but before digital I/Q demodulation). Using the chirp waveform generator, the performance of the Data Processor's pulse compression filter can be evaluated. Specifically, the most critical figure-of-merit is the range sidelobe level response of the processor, as these sidelobes affect how well sea surface clutter can be suppressed from the rain image.

The signal generator was programmed to simulate the chirp radar return for light rain scatterers at 1, 2, and 3 km altitudes above the ocean surface. The return power from the rain cells was set to -60 dB below the ocean return, which corresponds to the rain-to-surface return ratio for a 14 GHz TRMM-class radar viewing a 1 mm/h rain rate at nadir [4]. These return echoes were computed based on a chirp transmit pulse length of $51.2 \mu\text{s}$, a bandwidth of 4 MHz, and with a Hamming window taper to achieve the optimum match with the internal coefficients of the Data Processor's pulse-compression filter. The FPGA processors were also set to average the radar return samples in range by a factor of 2. At a 5 MHz baseband sampling rate for the processors, this range averaging translates into a vertical sampling interval of 60 m per pixel.

The resultant pulse-compressed, power detected output, plotted in Fig. 4, shows that the rain targets are detectable directly above the ocean even for very light rain rates of 1 mm/h. Note from the graph that the range sidelobe levels from the ocean reflection (the detected power between the rain scatterers) extend 75 dB below the ocean return. This range sidelobe performance in the FPGA processor enables the PR-2 to detect light rainfall




Figure 4: Pulse-compressed output recorded from the FPGA-based Data Processor during thermal chamber testing. (Light rain targets are located at 1, 2, and 3 km above the ocean surface.) Very low range sidelobes of -75 dB from the sea surface clutter indicate a high detection sensitivity for rain (down to a 1 mm/h rainfall rate).

condition	chassis net load	FPGA junction temperature T_J
1) One Data Processor configured and running at full throughput	28.1 W	35.1°C
2) No Data Processors configured, CTU configured and running	13.3 W	27.9°C
3) No FPGAs configured (quiescent load)	12.3 W	27.9°C

Table 2: Chassis power consumption and Data Processor FPGA temperature data at 25°C ambient temperature.

with a high level of sensitivity, even in the presence of high-energy clutter from the sea surface.

B. Power and Thermal Characteristics

During the free-run testing of the PR-2 processor, measurements of the chassis load currents and FPGA device temperatures were made under a variety of configuration conditions to infer the power consumption rating for each chip. Table 2 lists the power consumption data and FPGA temperature diode measurements at an ambient temperature $T_A = 25^\circ\text{C}$ for the cases where: 1) one Data Processor chip and the CTU are configured and running at the full throughput rate of 20×10^9 op/s; 2) neither Data Processor is configured, but the CTU is configured and running full sweep sequences; and 3) none of the Virtex FPGAs are configured.

From Table 2, the Data Processor power consumption can be calculated as the difference in net chassis powers between conditions (1) and (2) ($P_{DP} = 14.8$ W per FPGA chip). Likewise for the CTU chip, the power consumption is estimated at 1.0 W, i.e., the power difference between conditions (2) and (3). The junction temperature of the Data Processor FPGA is inferred from voltage measurements of an internal diode that has a temperature coefficient of -1.38 mV/°C. At full throughput and an air

Data Processor power consumption/temperature	$P_{DP} = 14.8 \text{ W per chip}$ $T_J = 35.1^\circ\text{C}$ (25 °C ambient)
Data Processor junction-to-air thermal resistance	$\theta_{JA} = 0.68^\circ\text{C/W}$ (mounted in chassis)
CTU power consumption	$P_{CTU} = 1.0 \text{ W}$
chassis quiescent power	$P_Q = 12.3 \text{ W}$
chassis net power budget	$2P_{DP} + P_{CTU} + P_Q = 42.9 \text{ W}$

Table 3: FPGA power and thermal summary for the PR-2 processor/controller boards.

temperature of 25°C, the Data Processor junction temperature stabilized to $T_J = 35.1^\circ\text{C}$.

With knowledge of the FPGA power consumption and the temperature differential between the die and ambient air, the chassis' effectiveness in removing heat from the FPGA chip can be evaluated using Ohm's Law for heat transfer,

$$T_J = T_A + \theta_{JA}P_{DP}, \quad (1)$$

where $\theta_{JA} = \theta_{JC} + \theta_{CA}$ is the junction-to-air thermal resistance, and θ_{JC} and θ_{CA} are the thermal resistances from junction-to-case and case-to-air. Solving for θ_{JA} in (1) yields a thermal resistance of 0.68°C/W . For comparison, the junction-to-case thermal resistance of the Virtex-1000 package (a 560-pin ball grid array) has a typical quoted rating of $\theta_{JC} = 0.8^\circ\text{C/W}$ —a value that is actually slightly higher than the empirical θ_{JA} . The implication is that the conduction of heat away from the FPGA is limited by the device package itself and not by the aluminum chassis, and therefore the machined heat sink mounts in the chassis design are highly efficient in removing the 15 W of heat away from each Data Processor chip.

In the thermal chamber tests it was also verified that the Data Processor could operate and produce valid pulse-compressed data over a -20 to $+70^\circ\text{C}$ ambient temperature range. For the worst-case tests at $T_A = 70^\circ\text{C}$ and free-running operation, the Data Processor die temperature stabilized at $T_J = 83.0^\circ\text{C}$, which is still below the 85°C design constraint for meeting the worst case timing requirements in a space-grade FPGA. Table 3 summarizes the power and thermal data results for the FPGA processor.

C. Error Detection and Correction

The design approach taken for error detection and correction (EDAC) of the volatile logic in the Virtex FPGAs is to scrub the configuration memory and periodically reconfigure all Virtex parts in the system. Such an approach has the advantage of clearing out single-event upsets (SEUs) without the added power consumption and mass of a triple-module redundancy circuit.

In the PR-2 design, the Xilinx FPGA configuration data and initial memory contents are loaded into a protected memory bank. Each memory word contains 16 bits of raw data plus an appended 6 bit Hamming code. A one-time programmable, antifuse-based Actel FPGA, which is immune to reconfiguration by SEUs, continuously reads the memory bank one word at a time and corrects any single-bit errors as they appear. After a nominal period of 5 minutes, the Actel part reads out the mem-

ory contents and reconfigures the Virtex FPGAs and their command tables so that the PR-2 processor can recover from these radiation-induced bit flips. In this way, the instrument cannot be permanently disabled by an SEU.

The reliability of the EDAC feature was tested in the thermal chamber by writing to the protected memory with one of the data bits stuck in a logic-low state. In the first period, the Virtex FPGAs would not configure with the corrupted data, but on the next cycle the memory contents were corrected and the FPGAs initialized successfully. This memory scrubbing test was repeated at several ambient temperature settings over range (25, 40, 50, 60, 70, and -20°C). In all cases, the FPGAs successfully recovered from the software-induced bit errors.

V. TRL DISCUSSION AND SUMMARY

The development of the PR-2's on-board processing hardware represents a breakthrough in realizing the observation capabilities for an advanced rain radar mission. Experiments in the lab with the prototype processor/controller boards has led to the following important findings: 1) the timing solution for auto-targeting of rain can be implemented in a single FPGA chip, and with a timing efficiency of 83–94% over the total available scan time; 2) pulse-compression of the chirp radar signal can be accomplished with very low range sidelobes (down to -75 dB) using a novel FPGA-based design for the receive processor; 3) the power consumption of each Data Processor FPGA is confirmed at 15 W at the full throughput rate, and the net power budget for the entire processor/controller is 43 W; 4) the conductively cooled design of a space-grade chassis yields a low, 10°C temperature gradient between the FPGA substrate and the ambient air temperature, thus improving the long-term reliability of the device; 5) the digital electronics successfully recovered from software-induced bit flips using a memory scrubbing technique, which points toward a solution to SEUs encountered in space. These results help establish the baseline instrument requirements and bring us one step closer toward a spaceborne mission for the PR-2.

Table 4 summarizes the increase in Technology Readiness Level for the design aspects of the PR-2 on-board processor. Many of these aspects have risen from concept formulation (TRL 2) to component or subsystem validation in a relevant environment (TRL 5–6). On-board data processing is now at the highest maturity level. The core components of this design—the functions for pulse-compression and averaging of multiple radar looks—have been proven in the lab with a target echo simulator as well as in several airborne precipitation experiments. The unique timing solution for the PR-2 began as a theoretical formulation for guaranteeing the auto-targeting of rain with collision-free transmit and receive pulse sequences, and has now matured into working FPGA firmware for the radar's Control and Timing Unit. This CTU firmware also represents a technical step forward in that it is the first high density (100,000 gate) FPGA-based timing controller designed for a spaceborne sensor. Tests of the CTU demonstrate the efficacy of using dedicated state-machine logic instead of a microprocessor for time-critical con-

Aspect of design	TRL at project start	Present TRL
On-board processing and science data compression	5	6
Timing solution for auto-targeting of rain	2	5
FPGA-based control of spaceborne instrument	2	5
Radiation upset correction approach	2	5

Table 4: TRL assessment for the PR-2 satellite processor/controller.

trol of the radar. Finally, a memory scrubbing approach for correcting radiation upsets in the FPGA configuration data was developed and proven as an attractive alternative to more complicated triple-module redundancy techniques.

ACKNOWLEDGMENT

The research described in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors thank Dr. Simone Tanelli for preparing the airborne PR-2 data products from the AMSR-E validation experiment.

REFERENCES

- [1] D. F. Everett, S. W. Bidwell, G. M. Flaming, T. B. Rykowski, E. F. Stocker, and L. E. Braatz, "Engineering of the Global Precipitation Measurement system," in *Proc. 2003 Aero. Conf.*, (Big Sky, MT), IEEE, Mar. 2003.
- [2] D. Atlas, "Adaptively pointing spaceborne radar for precipitation measurements," *J. Appl. Meteor.*, vol. 21, pp. 429–431, Mar. 1982.
- [3] A. C. Berkun, M. A. Fischman, and E. Im, "An advanced on-board processor and adaptive scanning controller for the Next-Generation Precipitation Radar," in *Proc. Earth Sci. Tech. Conf.*, (Pasadena, CA), NASA Earth Science Technology Office, June 2002.
- [4] R. Meneghini and T. Kozu, *Spaceborne Weather Radar*. Norwood, MA: Artech House, 1990.